

CLAIMS

- 1 1. A packet switching fabric comprising:
2 means forming a data ring, means forming a control ring, and means forming a plurality of data
3 communication network links each having at least one network node coupled thereto; and
4 a plurality of output queuing controlled switching devices coupled together by said data
5 ring means and said control ring means so that said network links can be selectively
6 communicatively coupled, each said switching device including,
7 data ring processing means for transmitting and receiving bursts of data to and
8 from adjacent ones of said devices via said data ring means,
9 network interface means having at least one network port for transmitting and
10 receiving data packets to and from said network links,
11 packet buffer means for storing the received data packets,
12 source managing means communicatively coupled to said data ring processing
13 means, and providing an interface between said network interface means and said packet
14 buffer means, said source managing means being operative to develop pointer
15 information for storing and reading each one of said received data packets to and from
16 said packet buffer means, and also being operative to develop destination identification
17 information associated with each one of said received data packets, said destination
18 identification information indicating an associated destination one of said network ports
19 of an associated destination one of said devices, said received data packets including
20 transfer packets received at an associated source one of said devices that is different from
21 said associated destination device, and local packets for which said associated destination
22 network port is a port of said associated source device, the pointer information and
23 destination identification information associated with each one of said local packets
24 providing local announcement information serving as a local request for access to said
25 associated destination network port, the pointer information and destination identification
26 information associated with each one of said transfer packets providing transfer
27 announcement information,
28 control ring processing means responsive to said transfer announcement
29 information, and operative to transmit and receive control messages to and from adjacent

30 ones of said devices via said control ring means, said control messages including
31 announcement messages each being associated with one of said transfer packets and
32 carrying said associated transfer announcement information, said control ring processing
33 means also being operative to provide received transfer announcement information in
34 response to each received one of said announcement messages, said received transfer
35 announcement information serving as a remote request for access to said associated
36 destination network port, and

37 destination managing means communicatively coupled to said network interface
38 means, and responsive to said local announcement information and said received transfer
39 announcement information, and operative to arbitrate between competing ones of said
40 local and remote requests for access to each of said network ports, and also operative to
41 generate transfer notification information associated with selected ones of said transfer
42 data packets;

43 said control messages further including transfer notification messages developed
44 by said control ring processing means in response to said transfer notification
45 information, each said notification message being associated with one of said transfer
46 data packets and carrying source identification information indicative of said associated
47 source device, each said notification message indicating to said associated source device
48 that said associated selected transfer data packet has been granted access to said
49 associated destination network port.

1 2. A packet switching fabric as recited in claim 1 wherein said transfer announcement
2 information associated with each one of said transfer data packets further includes source
3 identification information indicating said associated source device.

1 3. A packet switching fabric as recited in claim 1 wherein:

2 said network interface means includes a plurality of transmit buffer queues each
3 providing for receiving data from said destination managing means, and for transmitting bursts
4 of data to a corresponding one of said network links via a corresponding one of said network
5 ports; and

6 said destination managing means includes an output buffer manager for monitoring the
7 availability of buffer space in each of said transmit buffer queues, and wherein each of said
8 notification messages is transferred via said control ring means after a determination by said
9 output buffer manager that an associated destination one of said transmit buffer queues, that is
10 connected to said associated destination network port, includes a threshold amount of available
11 buffer space.

1 4. A packet switching fabric as recited in claim 3 wherein said output buffer manager is
2 operative to determine a number of blocks of buffer space available at each of said transmit
3 buffer queues, each of said available blocks providing buffer space sufficient for receiving a
4 burst of packet data from said destination managing means.

1 5. A packet switching fabric as recited in claim 4 wherein said notification information
2 further comprises an initial channel credit value indicating of a number of available blocks at the
3 destination transmit buffer queue associated with said selected transfer data packet prior to
4 transmitting said associated notification message.

1 6. A packet switching fabric as recited in claim 1 wherein said packet buffer means includes
2 at least one memory unit communicatively coupled with said source managing means via a
3 corresponding memory unit link, and wherein each of said devices receives a channel resource
4 patrol message from an adjacent one of said devices, said patrol message carrying channel
5 bandwidth information indicative of bandwidth available on said data ring means and bandwidth
6 available on said memory unit links, said control ring processing means being responsive to said
7 channel bandwidth information and operative to read and modify said channel bandwidth
8 information for the purpose of managing data transfer via said data ring means and via each of
9 said memory unit links.

1 7. A packet switching fabric as recited in claim 1 wherein:
2 said data ring means includes a plurality of data ring segments each coupling a corresponding
3 adjacent pair of said devices together;

4 said packet buffer means includes at least one memory unit communicatively coupled
5 with said source managing means via a corresponding memory unit link; and
6 each of said devices is responsive to a channel resource patrol message received from an
7 adjacent one of said devices, said patrol message carrying channel bandwidth information
8 including,
9 a plurality of data ring segment bandwidth parameters each being indicative of an
10 amount of bandwidth currently available at a corresponding one of said data ring
11 segments, and
12 a plurality of memory unit link bandwidth parameters each being indicative of an
13 amount of bandwidth currently available at a corresponding one of said memory unit
14 links;
15 said control ring processing means being responsive to said channel bandwidth
16 information, and being operative to read and modify said channel bandwidth information for the
17 purpose of managing data transfer via corresponding ones of a plurality of source-destination
18 channel paths for transmitting associated ones of said selected transfer data packets from said
19 associated source device to said associated destination device, each of said paths including
20 corresponding ones of said data ring segments and a corresponding one of said memory unit
21 links.

1 8. A packet switching fabric as recited in claim 7 wherein said patrol message is transferred
2 via said data ring, and wherein said control ring processing means is communicatively coupled
3 with said data ring processing means.

1 9. A packet switching fabric as recited in claim 7 wherein said control ring processing
2 means is operative to read selected sets of said data ring segment bandwidth parameters and said
3 memory unit link bandwidth parameters of said patrol message in response to said transfer
4 notification information, each of said selected sets of bandwidth parameters being associated
5 with one of said source-destination channel paths, said control ring processing means being
6 further operative to determine a maximum amount of bandwidth currently available for
7 transmitting data via each of said source-destination channel paths based on said associated
8 selected set of bandwidth parameters.

1 10. A packet switching fabric as recited in claim 9 wherein said control ring processing
2 means is further operative to determine an initial channel rate value associated with each one of
3 said selected transfer data packets, each said initial channel rate value indicating an initial
4 channel rate for transmitting bursts of said associated selected transfer data packet from said
5 packet buffer means of said associated source device to said associated destination device via
6 said associated source-destination channel path.

1 11. A packet switching fabric as recited in claim 10 wherein each of said notification
2 messages comprises:
3 a destination identification field for carrying said associated destination identification
4 information;
5 a source identification field for carrying said associated source identification information;
6 a packet location pointer field for carrying said associated pointer information;
7 an initial channel credit field for carrying said associated initial channel credit value; and
8 an initial channel rate field for carrying said associated initial channel rate value.

1 12. A packet switching fabric as recited in claim 11 wherein:
2 said control ring processing means is responsive to received ones of said notification messages,
3 and operative to provide received notification information associated with each one of said
4 received notification messages, said received notification information including said associated
5 source identification information, said associated destination identification information, said
6 associated pointer information, said associated initial channel credit value, and said associated
7 initial channel rate value; and
8 said source managing means is further responsive to said received notification
9 information, and operative to transfer data bursts of said associated selected transfer data packet
10 from said packet buffer means to said associated destination device via said associated source-
11 destination channel path in accordance with said associated initial credit value and said
12 associated initial channel rate value.

1 13. A packet switching fabric as recited in claim 12 wherein said source managing means
2 further comprises:
3 a source channel control unit responsive to said received notification information, and
4 operative to generate an initial channel data transfer signal associated with each one of said
5 received notification messages, each of said initial channel data transfer signals being repeatedly
6 activated a specified number of times in accordance with said associated initial channel rate
7 value, said specified number being determined based on said associated initial channel credit
8 value; and
9 a packet buffer control unit communicatively coupled to said packet buffer means and to
10 said network ports, said packet buffer control unit being responsive to said initial channel data
11 transfer signals, and being operative to read said specified number of data bursts of said
12 associated selected transfer data packet from said packet buffer means in accordance with said
13 associated initial channel rate value.

1 14. A packet switching fabric as recited in claim 13 wherein said source channel control unit
2 comprises:
3 a channel memory means responsive to said received notification information, and being
4 operative to manage channel information associated with corresponding ones of said source-
5 destination channel paths, said channel information having,
6 said associated source identification information, said associated pointer
7 information, a current channel credit value indicative of a number of bursts of said
8 selected data packet to be transmitted via said corresponding source-destination channel
9 path, said current channel credit value being initialized to said initial channel credit value,
10 and
11 a current channel rate value indicative of a channel rate for transmitting bursts of
12 said selected data packet via said corresponding source-destination channel path, said
13 current channel rate value being initialized to said initial channel rate value; and a
14 channel rate timer associated with said corresponding source-destination channel path,
15 said channel rate timer being responsive to said current channel rate value, and being
16 operative to generate a channel rate control signal that is repeatedly activated in
17 accordance with said current channel rate value.

1 15. A packet switching fabric as recited in claim 13 wherein said source channel control unit
2 includes:
3 a channel memory means responsive to said received notification information, and being
4 operative to manage channel information associated with said selected source-destination
5 channel path, said channel information including said associated source identification
6 information, said associated pointer information, said initial channel credit value, and said initial
7 channel rate value; and
8 a channel rate timer associated with said selected source-destination channel path, said
9 channel rate timer being responsive to said current channel rate value, and being operative to
10 generate said initial channel data transfer signal.

1 16. A packet switching fabric as recited in claim 13 wherein:
2 said output buffer manager is further operative to generate incremental credit transfer
3 information associated with said selected transfer data packet, said incremental credit transfer
4 information indicating an incremental number of available blocks at said destination transmit
5 buffer queue, said incremental number of available blocks having become available since said
6 transmission of said associated notification message;
7 said control messages further include incremental credit transfer messages developed by
8 said control ring processing means in response to said incremental credit transfer information,
9 said incremental credit transfer message carrying said incremental credit transfer information;
10 said control ring processing means is also responsive to terminated ones of said
11 incremental credit transfer messages, and operative to provide received incremental credit
12 transfer information; and
13 said source channel control unit is responsive to said received incremental credit transfer
14 information, and operative to increase said initial channel credit value.

1 17. A packet switching fabric as recited in claim 13 wherein:
2 said control ring processing means is further operative to generate incremental rate
3 transfer information associated with said selected transfer data packet, said incremental rate

4 transfer information indicating an incremental channel rate, said incremental channel rate being
5 determined based on additional channel bandwidth for transmitting data via said associated
6 source-destination channel path, said additional channel bandwidth being indicated by said patrol
7 message, said additional channel bandwidth having become available since said transmission of
8 said associated notification message;

9 said control messages further include incremental rate transfer messages developed by
10 said control ring processing means in response to said incremental rate transfer information, said
11 incremental rate transfer message carrying said incremental rate transfer information; and

12 said control ring processing means is also responsive to terminated ones of said
13 incremental rate transfer messages, and operative to provide received incremental rate transfer
14 information; and

15 said source channel control unit is responsive to said received incremental rate transfer
16 information, and operative to increase said channel rate.

1 18. A packet switching fabric as recited in claim 3 wherein each of said data packets is
2 received via a corresponding source port of said network ports, wherein each of said received
3 data packets includes header information specifying a corresponding destination address of a
4 corresponding destination node, and wherein said source managing means further includes a
5 packet routing control unit communicatively coupled to said network ports, and being responsive
6 to said destination addresses, and being operative to generate said destination information
7 associated with each one of said received data packets, said associated destination network port
8 being communicatively coupled to said corresponding destination node.

1 19. A packet switching fabric as recited in claim 18 wherein said packet routing control unit
2 is further operative to append each of said data bursts with block header information including
3 said destination ID value, and an end of packet indicator for indicating whether said data burst is
4 a last data burst of said data packet.

1 20. A packet switching fabric as recited in claim 1 wherein each of said control messages
2 includes a message field for indicating a message type of said control message

1 21. A packet switching fabric as recited in claim 19 wherein:
2 said network interface means further includes a receive buffer queue connected to each of
3 said network ports, each said receive buffer queue having an input connected to receive data
4 packets from a corresponding one of said network ports, and an output connected via a bus to
5 said packet routing control means and also to said packet buffer; and
6 said destination managing unit includes a data distribution control unit coupled to receive
7 said data bursts received by said data processing means, and having a plurality of outputs each
8 connected to one of said network ports via a corresponding one of a plurality of transmit buffer
9 queues, said data distribution control means including means for reading header information of
10 said data bursts and distributing said data bursts to said corresponding said destination ports.

1 22. A packet switching fabric as recited in claim 21 wherein said data distribution control
2 means includes a multicast queue for distributing multicast data bursts, having header
3 information specifying multicast addresses, to corresponding multiple ones of said transmit
4 buffer queues for transmission to multiple destination nodes.

1 23. A packet switching fabric as recited in claim 3 wherein said transmit queue buffers are
2 not large enough to store a whole one of said data packets and wherein cut through packet
3 transfer is implemented through said transmit buffer queues.

1 24. A packet switching fabric as recited in claim 6 wherein each of said memory units is
2 implemented by a dynamic RAM memory unit.

1 25. A packet switching fabric as recited in claim 1 wherein at least one of said network links
2 is an Ethernet link having a bandwidth of 100 Mbps.

1 26. A packet switching fabric as recited in claim 1 wherein at least one of said network links
2 is an Ethernet link having a bandwidth of 1 Gbps.

1 27. A packet switching fabric comprising:
2 means forming a data ring including a plurality of data ring segments;

3 means forming a control ring;

4 means forming a plurality of data communication network links each having at least one
5 network node coupled thereto; and

6 a plurality of switching devices coupled together by said data ring means and said control
7 ring means so that said network links can be selectively communicatively coupled, each adjacent
8 pair of said switching devices being coupled together by a corresponding one of said data ring
9 segments, each said switching device including,

10 network interface means having at least one network port for transmitting and
11 receiving data packets to and from said network links,

12 packet buffer means having at least one memory unit for storing the received data
13 packets,

14 data ring processing means for transmitting and receiving data bursts of said
15 received data packets to and from adjacent ones of said devices via said data ring means,

16 source managing means communicatively coupled to each of said memory units
17 via a corresponding memory unit link, and also communicatively coupled to said data
18 ring processing means and said network interface means, said source managing means
19 being operative to develop pointer information for storing and reading each one of said
20 received data packets to and from said packet buffer means, and also being operative to
21 develop destination identification information associated with each one of said received
22 data packets, said destination identification information indicating an associated
23 destination one of said network ports of an associated destination one of said devices,

24 control ring processing means operative to develop, transmit and receive control
25 messages to and from adjacent ones of said devices via said control ring means, said
26 control messages for controlling packet transfer operations including transmitting
27 associated selected ones of said received data packets from said associated source device
28 to said associated destination device via an associated source-destination channel path
29 including associated ones of said data ring segments and an associated one of said
30 memory unit links, said control ring processing means being responsive to a channel
31 resource patrol message received from an adjacent one of said devices, said patrol
32 message carrying channel bandwidth information indicative of bandwidth available on
33 said data ring means and bandwidth available on said memory unit links, said control ring

34 processing means being responsive to said channel bandwidth information and operative
35 to read and modify said channel bandwidth information for the purpose of managing data
36 transfer via said data ring means and via each of said memory unit links, and
37 destination managing means for receiving data bursts of said received data
38 packets from said data ring processing means, and for providing said data bursts to said
39 associated destination network ports.

1 28. A packet switching fabric as recited in claim 27 wherein:

2 said received data packets include transfer packets received at an associated source one of
3 said devices that is different from said associated destination device, and local packets for which
4 said associated destination network port is a port of said associated source device, the pointer
5 information and destination identification information associated with each one of said local
6 packets providing local announcement information serving as a local request for access to said
7 associated destination network port, the pointer information and destination identification
8 information associated with each one of said transfer packets providing transfer announcement
9 information;

10 said control ring processing means is responsive to said transfer announcement
11 information, said control messages including announcement messages each being associated
12 with one of said transfer data packets and carrying said associated transfer announcement
13 information, said control ring processing means also being operative to provide received transfer
14 announcement information in response to each received one of said announcement messages,
15 said received transfer announcement information serving as a remote request for access to said
16 associated destination network port;

17 said destination managing means is responsive to said local announcement information
18 and said received transfer announcement information, and operative to arbitrate between
19 competing ones of said local and remote requests for access to each of said network ports, and
20 also operative to generate transfer notification information associated with selected ones of said
21 transfer data packets; and

22 said control messages further include transfer notification messages developed by said
23 control ring processing means in response to said transfer notification information, each said
24 notification message being associated with one of said transfer data packets and carrying source

25 identification information indicative of said associated source device, each said notification
26 message indicating to said associated source device that said associated selected transfer data
27 packet has been granted access to said associated destination network port.

1 29. A packet switching fabric as recited in claim 27 wherein said channel bandwidth
2 information comprises:
3 a plurality of data ring segment bandwidth parameters each being indicative of an amount
4 of bandwidth currently available at a corresponding one of said data ring segments; and
5 a plurality of memory unit link bandwidth parameters each being indicative of an amount
6 of bandwidth currently available at a corresponding one of said memory unit links.

1 30. A packet switching fabric as recited in claim 28 wherein said transfer announcement
2 information associated with each one of said transfer data packets further includes source
3 identification information indicating said associated source device.

1 31. A packet switching fabric as recited in claim 28 wherein:
2 said network interface means includes a plurality of transmit buffer queues each
3 providing for receiving data from said destination managing means, and for transmitting bursts
4 of data to a corresponding one of said network links via a corresponding one of said network
5 ports; and
6 said destination managing means includes an output buffer manager for monitoring the
7 availability of buffer space in each of said transmit buffer queues, and wherein each of said
8 notification messages is transferred via said control ring means after a determination by said
9 output buffer manager that an associated destination one of said transmit buffer queues, that is
10 connected to said associated destination network port, includes a threshold amount of available
11 buffer space.

1 32. A packet switching fabric as recited in claim 31 wherein said output buffer manager is
2 operative to determine a number of blocks of buffer space available at each of said transmit
3 buffer queues, each of said available blocks providing buffer space sufficient for receiving a
4 burst of packet data from said destination managing means.

1 33. A packet switching fabric as recited in claim 32 wherein said notification information
2 further comprises an initial channel credit value indicating of a number of available blocks at the
3 destination transmit buffer queue associated with said selected transfer data packet prior to
4 transmitting said associated notification message.

1 34. A packet switching fabric as recited in claim 33 wherein said patrol message is
2 transferred via said data ring, and wherein said control ring processing means is
3 communicatively coupled with said data ring processing means.

1 35. A packet switching fabric as recited in claim 34 wherein said control ring processing
2 means is operative to read selected sets of said data ring segment bandwidth parameters and said
3 memory unit link bandwidth parameters of said patrol message in response to said transfer
4 notification information, each of said selected sets of bandwidth parameters being associated
5 with one of said source-destination channel paths, said control ring processing means being
6 further operative to determine a maximum amount of bandwidth currently available for
7 transmitting data via each of said source-destination channel paths based on said associated
8 selected set of bandwidth parameters.

1 36. A packet switching fabric as recited in claim 35 wherein said control ring processing
2 means is further operative to determine an initial channel rate value associated with each one of
3 said selected transfer data packets, each said initial channel rate value indicating an initial
4 channel rate for transmitting bursts of said associated selected transfer data packet from said
5 packet buffer means of said associated source device to said associated destination device via
6 said associated source-destination channel path.

1 37. A packet switching fabric as recited in claim 36 wherein each of said notification
2 messages comprises:
3 a destination identification field for carrying said associated destination identification
4 information;
5 a source identification field for carrying said associated source identification information;

6 a packet location pointer field for carrying said associated pointer information;
7 an initial channel credit field for carrying said associated initial channel credit value; and
8 an initial channel rate field for carrying said associated initial channel rate value.

1 38. A packet switching fabric as recited in claim 37 wherein:

2 said control ring processing means is responsive to received ones of said notification
3 messages, and operative to provide received notification information associated with each one of
4 said received notification messages, said received notification information including said
5 associated source identification information, said associated destination identification
6 information, said associated pointer information, said associated initial channel credit value, and
7 said associated initial channel rate value; and

8 said source managing means is further responsive to said received notification
9 information, and operative to transfer data bursts of said associated selected transfer data packet
10 from said packet buffer means to said associated destination device via said associated source-
11 destination channel path in accordance with said associated initial credit value and said
12 associated initial channel rate value.

1 39. A packet switching fabric as recited in claim 38 wherein said source managing means
2 further comprises:

3 a source channel control unit responsive to said received notification information, and
4 operative to generate an initial channel data transfer signal associated with each one of said
5 received notification messages, each of said initial channel data transfer signals being repeatedly
6 activated a specified number of times in accordance with said associated initial channel rate
7 value, said specified number being determined based on said associated initial channel credit
8 value; and

9 a packet buffer control unit communicatively coupled to said packet buffer means and to
10 said network ports, said packet buffer control unit being responsive to said initial channel data
11 transfer signals, and being operative to read said specified number of data bursts of said
12 associated selected transfer data packet from said packet buffer means in accordance with said
13 associated initial channel rate value.

1 40. A packet switching fabric as recited in claim 39 wherein said source channel control unit
2 comprises:

3 a channel memory means responsive to said received notification information, and being
4 operative to manage channel information associated with corresponding ones of said source-
5 destination channel paths, said channel information having,

6 said associated source identification information, said associated pointer
7 information, a current channel credit value indicative of a number of bursts of said
8 selected data packet to be transmitted via said corresponding source-destination channel
9 path, said current channel credit value being initialized to said initial channel credit value,
10 and

11 a current channel rate value indicative of a channel rate for transmitting bursts of
12 said selected data packet via said corresponding source-destination channel path, said
13 current channel rate value being initialized to said initial channel rate value; and a
14 channel rate timer associated with said corresponding source-destination channel path,
15 said channel rate timer being responsive to said current channel rate value, and being
16 operative to generate a channel rate control signal that is repeatedly activated in
17 accordance with said current channel rate value.

1 41. A packet switching fabric as recited in claim 39 wherein said source channel control unit
2 includes:

3 a channel memory means responsive to said received notification information, and being
4 operative to manage channel information associated with said selected source-destination
5 channel path, said channel information including said associated source identification
6 information, said associated pointer information, said initial channel credit value, and said initial
7 channel rate value; and

8 a channel rate timer associated with said selected source-destination channel path, said
9 channel rate timer being responsive to said current channel rate value, and being operative to
10 generate said initial channel data transfer signal.

1 42. A packet switching fabric as recited in claim 39 wherein:

2 said output buffer manager is further operative to generate incremental credit transfer
3 information associated with said selected transfer data packet, said incremental credit transfer
4 information indicating an incremental number of available blocks at said destination transmit
5 buffer queue, said incremental number of available blocks having become available since said
6 transmission of said associated notification message;

7 said control messages further include incremental credit transfer messages developed by
8 said control ring processing means in response to said incremental credit transfer information,
9 said incremental credit transfer message carrying said incremental credit transfer
10 information;

11 said control ring processing means is also responsive to terminated ones of said
12 incremental credit transfer messages, and operative to provide received incremental credit
13 transfer information; and

14 said source channel control unit is responsive to said received incremental credit transfer
15 information, and operative to increase said initial channel credit value.

1 43. A packet switching fabric as recited in claim 39 wherein:

2 said control ring processing means is further operative to generate incremental rate
3 transfer information associated with said selected transfer data packet, said incremental rate
4 transfer information indicating an incremental channel rate, said incremental channel rate being
5 determined based on additional channel bandwidth for transmitting data via said associated
6 source-destination channel path, said additional channel bandwidth being indicated by said patrol
7 message, said additional channel bandwidth having become available since said transmission of
8 said associated notification message;

9 said control messages further include incremental rate transfer messages developed by
10 said control ring processing means in response to said incremental rate transfer information, said
11 incremental rate transfer message carrying said incremental rate transfer information; and

12 said control ring processing means is also responsive to terminated ones of said
13 incremental rate transfer messages, and operative to provide received incremental rate transfer
14 information; and

15 said source channel control unit is responsive to said received incremental rate transfer
16 information, and operative to increase said channel rate.